Evaluation

Online learning

links with optimization and games

Université Paris–Saclay

convex optimization with line-search

Let $\mathcal{X} \subset \mathbb{R}^d$ be closed convex set, $f : \mathbb{R}^d \to \mathbb{R}$ a differentiable convex function that admits a minimizer x_* on $\mathcal{X}$, meaning:

$$f(x_*) = \min_{x \in \mathcal{X}} f(x).$$

Let $H : \mathbb{R}^d \to \mathbb{R} \cup \{+\infty\}$ be a mirror map compatible with $\mathcal{X}$, $(\gamma_t)_{t \geqslant 1}$ a positive sequence and $((x_t, y_t))_{t \geqslant 1}$ be such that it is a strict UMD sequence associated with regularizer $h = H + I_{\mathcal{X}}$ and sequence $(-\gamma_t \nabla f(x_t))_{t \geqslant 1}$.

1) Explicitely write at least two different sequences $(x_t)_{t \geqslant 1}$ that satisfy the above.

2) Assume that for all $t \geqslant 1$,

$$\gamma_t D_f(x_{t+1}, x_t) \leqslant D_h(x_{t+1}, x_t;\ y_t) \qquad (1)$$

Establish a guarantee on

$$f(x_{T+1}) - f(x_*).$$

Hint: Get some inspiration from Proposition 6.2.1 from the lecture notes.

3) Let $\|\cdot\|$ be a norm and $K, L > 0$. We assume in this question that for $\|\cdot\|$, f is L-smooth and h is K-strongly convex. Then derive a guarantee for a simple choice of $(\gamma_t)_{t\geqslant 1}$.

4) For a given $x \in \mathscr{X}$ and $y \in \partial h(x)$, does a step-size $\gamma > 0$ satisfying

$$x' = \nabla h^*(y - \gamma \nabla f(x)) \quad \text{and} \quad \gamma D_f(x', x) \leqslant D_h(x', x;\ y)$$

always exist?

5) Propose a practical method for finding, given x_t and y_t, a step-size γ_t satisfying condition (1).

6) Now consider for H the Euclidean mirror map. Perform numerical experiments for two different convex and differentiable objective functions (e.g. logistic/linear regression). Compare the case $\mathscr{X} = \mathbb{R}^d$ and the case where $\mathscr{X}$ is a closed Euclidean ball. In the latter case, compare the two different algorithms written in question 1).